arXiv:1303.2595v1 [cs.DB] 11 Mar 2013

Integrating Space, Time, Version and Scale Using Alexandrov Topologies

Norbert Paul, Patrick Erik Bradley, and Martin Breunig

Karlsruhe Institute of Technology, 76128 Karlsruhe, Germany

Abstract. As a contribution to higher dimensional spatial data modelling this article introduces a novel approach to spatial database design. Instead of extending the canonical Solid-Face-Edge-Vertex schema of topological data, these classes are replaced altogether by a common type `SpatialEntity`, and the individual "bounded-by" relations between two consecutive classes are replaced by one separate binary relation `BoundedBy` on `SpatialEntity`. That relation defines a so-called *Alexandrov topology* on `SpatialEntity` and thus exposes the fundamental mathematical principles of spatial data design. This has important consequences: First, a formal definition of topological "dimension" for spatial data can be given. Second, every topology for data of arbitrary dimension has such a simple representation. Also, version histories have a canonical Alexandrov topology, and generalisations can be consistently modelled by the new consistency rule *continuous functions* between LoDs, and *monotonicity* enables accelerated path queries. The result is a relational database schema for spatial data of dimension 6 and more which seamlessly integrates 4D space-time, levels of details and version history. Topological constructions enable queries across these different aspects.

1 Introduction

2D and 3D spatial models are well established for spatial data modelling, and there exist standards like CityGML in geo-spatial modelling and IFC for architectural models. Currently there is active research on spatio-temporal queries [1] as well as 4D spatio-temporal modelling [2], and also on considering other aspects like scale [3] as additional dimensions of spatial data. Also integrating the versioning graph, by itself 1D, pushes up the dimension upper bound.

Besides that, research on nD modelling provides generic spatial data models without a fixed dimension upper bound [4] and often gives a formal definition of "topological dimension" of spatial data. Topology has its own sub-discipline called "dimension theory" [5] where the possible definitions of such "topological dimension" are investigated. Among these, the *Krull dimension* [6, p. 5] is particularly applicable for topological data and is proposed here as a standard definition of spatial data dimension. Throughout this article "dimension" of spatial data is the Krull dimension of the topological space established by the data entities.

Challenged by a scientist's request to show that there are applications for data of dimension beyond 4, and inspired by VAN OOSTEROM's ideas [7], this article demonstrates the mathematical foundations of generic nD spatial modelling and its use for combining 3D spatial data, time, scale, and versioning into an integrated 6D+ model. In particular, it will be shown that sensible integration of scale increases the dimension of spatial data by more than one, hence the "+" in 6D+. Even more, further increasing the dimension can *decrease* the complexity of the data model, which renders statements like "nD modelling leads to a combinatorial explosion of complexity" wrong.

The reader is assumed familiar with basic concepts of mathematical topology. To facilitate the lecture, references are given to locations in textbooks explaining the applied concepts in more detail. As the main intent is to expose the mathematics of nD modelling, the relational model, formulated by CODD [8], is used. Its sound mathematical basis alleviates its topologising in an elegant way [9], and the principles described here are also applicable to object oriented modelling. They also serve as a basis for a generic topological database model as an extension of the relational model: instead of tables, there are spaces on which queries operate and which they return as a result. The first author of this article has implemented an experimental prototype of such a topological relational database model which runs on `http://pavel.gik.kit.edu`.

2 Dimension

In 3D spatial data usually four kinds of topological entities are represented in computer storage: *Vertices* are zero-dimensional discrete points in $\mathbb{R}^3$. *Edges* are one-dimensional manifolds which are the interior of paths starting at one vertex and ending at another vertex. In general, an edge is *bounded* by two vertices. *Faces* are two-dimensional manifolds enclosed by at least one loop of edges. Some models allow to specify more loops where one loop is the outer boundary and each additional loop is a hole boundary. Topologically, there is no difference between an "outer" boundary loop and a "hole" boundary loop because every hole can be made an outer boundary by stretching it wide enough and then flipping the face over. To wrap this up, a face is bounded by edges. *Solids* are three-dimensional manifolds enclosed by a set of faces which constitute a cavity within which the solid resides. Such cavity is often called a *shell* and, again, a solid is bounded by faces. Most 3D models establish a "chain" of four classes with two consecutive classes connected by a "bounded-by" association (cf. [10] for an overview). Note that the chain length equals the model dimension.

According to 3D spatial models, time can be modelled by the real line $\mathbb{R}$. Each moment in time can be represented by a real number, thus resembling a vertex in $\mathbb{R}^3$. A time span is an open interval (t_1, t_2) bounded by a starting point t_1 and an ending point t_2, thus resembling an edge. This gives two classes of temporal entities: *Moment* and *Timespan* where each one-dimensional timespan is bounded by two zero-dimensional moments. According to spatial data

this "bounded by" association can be considered a special case of an association chain of length 1 which, again, corresponds with the dimension.

When data is changed over time several versions of the data exist. In versioning software two consecutive versions along a version history are commonly connected by an edge. Versions can also fork and merge and so a version history is a directed acyclic graph (DAG) with two classes, a *Version* and a *Transition* and each transition is bounded by an initial and a terminal version. So we could say that the "dimension" of a version history is one: there is only one association from *Transition* to *Version*. We will later take an alternative view on the version history which results in a different dimension.

Spatial data is often organised hierarchically at different levels of "scale" or "detail" (LoD). For smooth transitions between LoDs it is often proposed to interpolate between consecutive LoDs by continuous morphing [7]. Then the space between two consecutive LoDs is considered an edge bounded by two LoDs. So, we have a one-dimensional space.

Thus we have four types of spaces which we might call "elementary spaces": The "spatial" space which is the Euclidean 3D, the one-dimensional Euclidean temporal space, a one-dimensional version space, and another one-dimensional scale space which is essentially a linear graph $\cdots \to \bullet \to \bullet \to \cdots$.

Now a combination of these elementary spaces can be used to get higher-dimensional "combined spaces". In such a combined space a 3D *Solid* s and a 1D *Timespan* ts can be combined to a 4D pair (s, ts) which represents a 4D entity in space-time: the trajectory of s during ts. But for a zero-dimensional *Moment* in time t, the pair (s, t) is a 3D element in space-time representing that solid s at the moment t. When ts is bounded by t then the 4D-entity (s, ts) is bounded by a 3D-entity (s, t) and for each lower-dimensional nD element x in the bounded-by association chain of s there is a pair (x, t) in a corresponding association chain of (s, t). This means that in the 4D model the length of the association chains has increased by 1. Within that model each pair (a, b) satisfies the dimension formula $\dim(a, b) = \dim a + \dim b$.

As every database management system (DBMS) can only model finite sets we need a definition of "dimension" for finite spaces. As seen above, our nD spaces consist of entities, possibly distributed over several classes, and a bounded-by relation of chain length n. Hence "dimension" of spatial data is the maximal length of a chain of entities such that each is bounded by a consecutive element. We call this the *combinatorial dimension* of spatial data. The note [11] proves that this combinatorial dimension is equivalent to the topological Krull dimension [6, p. 5]. Note that even a simple sequence $(n, n-1, \ldots, 0)$ can be considered a set $\{n, n-1, \ldots, 0\}$ with "association" $a\,S\,b \Leftrightarrow a - 1 = b$ which then has a combinatorial dimension of n and obviously does not suffer from any "combinatorial explosion of complexity" whatsoever. We can even decrease the complexity of spatial data by increasing its dimension.

3 Spatial Dimension and Consistency

Usually, in a DBMS consistency rules are design tools, and it lies at the discretion of the user to make use of them. However, in spatial data modelling, "topological consistency" is often mentioned, but the rules to tell "consistent" from "inconsistent" spaces vary in the literature. We see two reasons: First, as mentioned above, the user should be entitled to decide which spatial data he considers "consistent" and which he does not. Second, the term "topological (in)consistency" is unknown in mathematics. Topology provides a rich set of well-defined topological *properties* which may or may not be used for a particular application as a consistency rule. We will discuss here some "consistency" for nD spatial entities and present applications where these rules do not apply.

A *vertex* is usually composed of an id, the coordinates, and additional attributes. The consistency rule is that a vertex is not bounded by another object. However, this property can be used to *define* vertex: In a topological space a *vertex* is an element that is not bounded by another element. A practical application is the room connectivity graph of a building: In this graph the rooms are the 0D vertices, and the doors are their connecting 1D edges. Each door is then "bounded by" its two rooms. This possibility of "shifting down" the dimension and inverting it by "flipping" the boundary relation is a characteristic of the topological spaces occurring in spatial data models. That "flipping" is intimately related to the Poincaré duality with which it should not be confounded.

Some spatial data models define faces by a cycle of vertices around that face. This gives implicit edges between two consecutive vertices in the cycle. But we will only consider models with explicitly given edges. Then an edge usually has two references to two boundary vertices. When we weaken that rule of exactly two vertices we can *define* "edge" as a spatial entity with non-empty boundary that consists only of vertices. This allows selection results where the selection predicate only applies to, say, two edges but only one common connecting vertex. It would then be unwise to forget the connectivity information in the selection result only because it fails to satisfy the global constraints. Moreover, if an edge is allowed to have more than one vertex, then the data model allows hypergraphs, an abstract topological structure often used in practice. Additionally, an edge with four vertices might be considered an edge with a hole that occurs e.g. in intersections with non-convex faces.

The classical consistency rule for faces is that the boundary edges must form loops. Mathematically, this means that a boundary must form a cycle, which is the fundamental property of a *chain complex* from algebraic topology [12]. However, it makes sense in general to permit query results violating consistency rules that have been set up for the underlying query input spaces.

Another consistency rule is hard-coded into the classical sequential class schema: The only possible association between a face and one of its boundary vertices is indirectly via an intermediate edge. The schema does not allow a direct face-vertex association that "bypasses" the edge class. However, a vertex within a face may be considered, say, a collapsed interior face, e.g. a city in a region at a lower LoD. If we now specify a superclass *SpatialEntity* of Solid, Face,

Edge, and Vertex and replace the associations between two consecutive classes by one association of that class to itself, that rule can be made optional, too.

Making the above consistency constraints optional immediately leads to a directed acyclic graph (DAG) of spatial entities. A DAG defines a partial ordering on its elements, and since 1937 it is well-known that partial orderings are essentially the same as the so-called T_0 *Alexandrov topologies* [13], or, to put it short, *spatial data models are topological spaces*. We will now provide an initial relational database schema for 3D spatial data based on this observation:

$$X(\underline{id}, attributes), R(\underline{ida}, \underline{idb}), \quad (ida) \xrightarrow{FK} X,\ (idb) \xrightarrow{FK} X \tag{1}$$
$$Vertex(\underline{vid}, x : \mathbb{R}, y : \mathbb{R}, z : \mathbb{R}), \quad (vid) \xrightarrow{\text{FK}} X\ .$$

Table X contains the spatial entities, and table R specifies the "bounded-by" relation. Primary key attributes are underlined. Attribute *attributes* is a placeholder for the set of "semantic", i.e. non-spatial attributes. The only consistency rule here is that R should be acyclic. Although every relation R defines an Alexandrov topology $\mathcal{T}(R)$ of the space

$$(X, \mathcal{T}(R)), \quad \mathcal{T}(R) := \{A \subseteq X \mid \forall (a, b) \in R : b \in A \Rightarrow a \in A\}$$

that space is T_0-separable iff R is acyclic [14, Ex. 133]. T_0-separability is a sensible consistency rule for spatial data, and will be maintained here. Note, however, that non-separable spaces also are valid topological spaces, and so from the mathematical viewpoint acyclicity is an optional consistency rule, too. For each topology-defining relation R on X we assume the existence of a view on its pre-order R^* which is the transitive and reflexive closure of R:

```
create view poR as
with recursive Rst(ida,idb) as (
   select id as ida, id as idb from X
   union
   select Rst.ida, R.idb
   from Rst join R on (Rst.idb = R.ida))
select * from Rst;
```

This schema allows spatial data of arbitrary dimension, but for the moment we assume that R chains have length ≤ 3, so the topological dimension matches the "geometric dimension", i.e. the number of coordinate attributes of *Vertex*. Also, the *id* of all vertices should occur in table *Vertex*. For obvious reasons, we call a pair (X, R) consisting of a set X and a binary relation R on X a *topological data type*, or sometimes simply *space*.

4 Temporal Dimension

Here, we establish the 4D *space-time* which models changes of 3D space over time. For illustrative reasons, a 2D space-time of a 1D "building" example in Lineland [15, ch. 13] changing over 1D time will be discussed first. Assume in

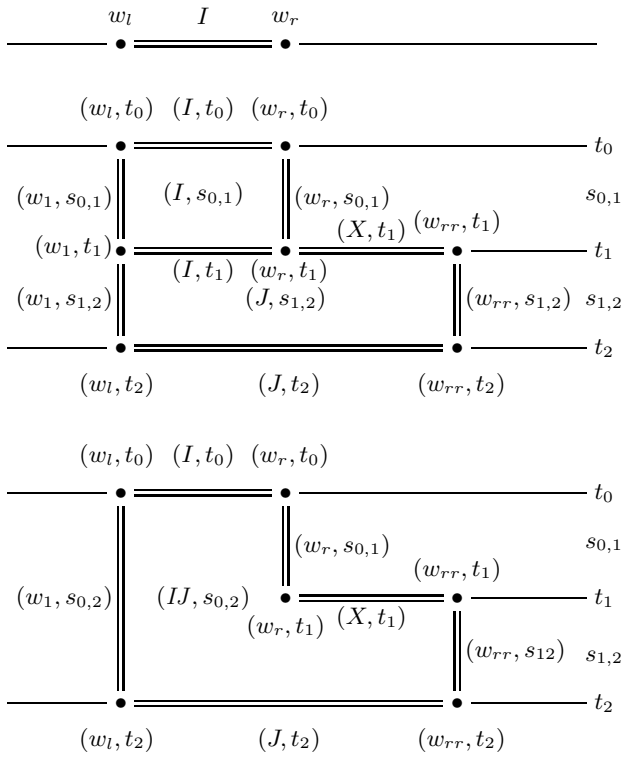

Fig. 1. The process of a 1D house in Lineland constructed at time t_0, extended by a portion X at time t_1, and demolished at time t_2, is modelled as a 2D space-time complex in two steps. The middle complex repeats each element at a given point or period of time, and the lower complex identifies some elements, the left wall w_l and some elements of the interior I, to get fewer entities to store.

1D Lineland a "house" with interior I and boundary walls w_l to the left and w_r to the right as depicted on top of Fig. 1. The boundary of I are the vertices w_l and w_r. The house is erected at time t_0, modelled by "tagging" I, w_l and w_r with t_0 giving the pairs (I, t_0), (w_l, t_0), and (w_r, t_0). The spatial "bounded-by" associations from I to w_l and from I to w_r are carried over as tagged space-time associations from (I, t_0) to (w_l, t_0) and from (I, t_0) to (w_r, t_0). The interior I, existing over a time-span $s_{0,1}$, is also tagged, giving its trajectory $(I, s_{0,1})$. This element is bounded in the horizontal (i.e. "spatial") direction by $(w_1, s_{0,1})$ because I is bounded by w_1. Here the "tag" $s_{0,1}$ does not change in this boundary association. The element $(I, s_{0,1})$ is also bounded in the vertical (i.e. "temporal") direction by (I, t_0). Here the boundary association fixes element I and is taken from the boundary association between $s_{0,1}$ and t_0.

The space at time point t_1 models the before-after change derived by a 1D overlay of the spatial model of I before, and the spatial model of J after the

change. Each overlay entity has a reference to its corresponding entity of the input spaces. Mathematically, these are two topologically continuous partial functions p (like "past") and f (like "future") from the overlay space back to the two input spaces. Such functions are called *attaching maps*. Now we can topologically paste [16, Ch. 3 §7] the overlay onto the "past" entities by specifying that a tagged image element $(I, s_{0,1})$ is bounded by an element (x, t_1) if the reference $p(x) = I$ holds. We can also do that attachment onto the "future" by specifying that $(w_1, s_{0,1})$ is bounded by (w_1, t_1) because $p(w_1) = w_1$. Interestingly (w_r, t_1) is a boundary element of the past wall trajectory $(w_r, s_{0,1})$ because of $p(w_r) = w_r$, but it is a boundary element of the future interior trajectory $(J, s_{1,2})$ because of $f(w_r) = J$.

This step is only an intermediate formal step which would create a lot of redundant data when implemented explicitly. To avoid redundancy a sequence of a tagged spatial entity not changing over time can be collapsed into one. Additionally, the user may specify further entities that, though having changed over time, are considered "identical" before and after the change. In our example this identification is carried out on $(w_1, s_{0,1})$, (w_1, t_1), and $(w_1, s_{1,2})$, and it collapses $(I, s_{0,1})$, (I, t_1), and $(J, s_{1,2})$ to $(IJ, s_{01,2})$. Mathematically, this gives a topological "quotient space" [16, Ch. 3] depicted in the lower part of Fig. 1.

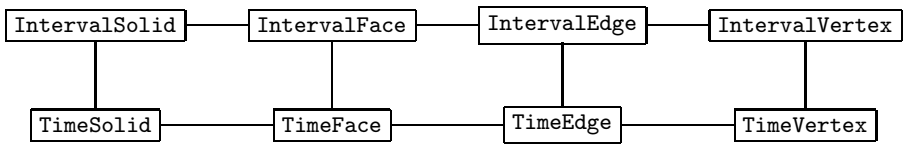

Fig. 2. Directly combining the classical sequential space model with the same classical approach for a temporal model creates eight classes for time-space entities.

The same construction applies to 3D spatial models to get a 4D space-time which is simply another topological space. However, when the 3D model is represented by a chain of three associations, and when the temporal model also has that classical layout, we get a "grid" of eight different classes of pairs of spatial-temporal entities (cf. Fig. 2). This results in ten different space-time-boundary associations, represented by lines between classes. On the other hand, if we had only two classes `SpatialEntity` and `TemporalEntity`, each with a relation `BoundedBy`, then there would be only one spatio-temporal class where each entity consists of a pair (`spatial`,`temporal`), and only one boundary relation associating each pair `(spatial,temporal)` with all (*ds*,`temporal`) for every boundary element *ds* of `spatial` and, dually, with the pairs `(spatial,`*dt*`)` for every boundary time point *dt* of `temporal`. This yields the topology of the so-called *product space* [16, Ch. I §3]. The class `SpaceTimeEntity` with a relation `BoundedBy` (cf. Fig. 3) allows to model arbitrary space-time configurations. When we consider *id* a surrogate key for pairs of space-time-entities the relational schema (1) for

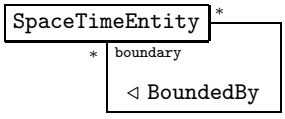

Fig. 3. The one space-time class and bounded-by association obtained by combining the spatial and temporal models in a non-classical way.

3D spatial data can be easily extended to cover spatio-temporal data:

$$\mathit{Vertex}(\underline{\mathit{vid}}, x : \mathbb{R}, y : \mathbb{R}, z : \mathbb{R}, t : \mathbb{R}) \ .$$

Note the simplicity. Like [1], it allows moving objects between two time points that may be geometrically interpolated by a query. A simple topological SQL-query for the space at a given time point t with no interpolation is:

```
create view Xt as
with mmX(id, tmin, tmax) as (
   select X.id, min(V.t) as tmin, max(V.t) as tmax
   from   (X join poR on (X.id = poR.ida))
          join Vertex V on (poR.idb = V.vid)
   group by X.id)
select X.* from X join mmX on (X.id = mmX.id)
where (mmX.tmin < t and t < mmX.tmax)
   or (mmX.tmin = t and t = mmX.tmax);
```

The sub-query mmX first associates each element $X.id$ with its topological closure $\mathrm{cl}(X.id)$ by joining it to poR. The join with *Vertex* selects all vertices of $\mathrm{cl}(X.id)$. The group-by clause then computes the time interval for $X.id$.

This query selects a sub*set* Xt of X. As the pair (X, R) defines a topological space $(X, \mathcal{T}(R))$, there should be a relation Rt that generates the sub*space* topology $\mathcal{T}(R)|_{Xt}$ (cf. [16, Ch. I §3]). Simply restricting R to Xt is generally wrong [9]. Naively restricting poR to Xt is correct but expensive. An optimal Rt is achieved by passing that restriction to Codd's `OPEN` operator which returns a minimal relation whose transitive closure is the same as that of the input relation [17, p. 427]. Hence $Rt := \mathtt{OPEN}(\{(a, b) \in poR \mid a \neq b, a \in Xt, b \in Xt\})$.

Just as the set Xt is a Θ-selection of X, the space (Xt, Rt) can be considered a topological Θ-selection of (X, R), hence a "topologised" basic relational query operator. The topology $\mathcal{T}(Rt)$ is the minimal topology for which the inclusion function $i : Xt \to X$ is continuous [14, Ex. 1]. A relationally complete topological database query language is given in [9]. A first simple prototype implementation can be found under `http://pavel.gik.kit.edu/`.

It is also noteworthy that an edge in X, representing the trajectory of a vertex v from t_0 to t_1 with $t_0 < t < t_1$, topologically becomes a vertex in the result space Xt which does not contain the edge's space-time-vertices.

5 Versioning

The *version graph* is a DAG whose vertices resp. edges correspond to the different versions resp. modifications of a space. An edge (i, j) indicates that version i has been modified to obtain version j. Clearly, not all versions need to be stored explicitly. It suffices to store the initial version, and any other version v can be reconstructed from it by applying all modifications on the paths to v in the version graph. It is also advisable to redunantly store more versions e.g the current version but this leads to the problem of balancing redundancy avoidance against robustness and speed and we will not delve into this matter. Note that we only consider topological changes here.

The possible topological modifications of a pair (X, R) are: First, a point can be added, and a point $x \in X$ can be removed. The latter forces relation R to be modified by removing all pairs (y, x) and (x, z) with $y, z \in X$. But if y, z are different from x with $(y, x) \in R$ and $(x, z) \in R$, then the pair (x, z) must be added to R if it is not already contained in R. The reason is that the modified set $X \setminus \{x\}$ should be a *subspace* of (X, R), meaning that the (possibly indirect) bounded-by association $y \to z$ in X must be retained in $X \setminus \{x\}$ by the induced bounded-by relation. Further, a pair of points (either new or not removed from X) can be added to R, and a pair can be removed from R. These are the elementary modifications, and a general modification is a sequence of elementary modifications.

The following relation schema:

$$X(\underline{id}, version, atts),\ version \xrightarrow{\text{FK}} VX,\ R(\underline{ida}, \underline{idb}),\ ida \xrightarrow{\text{FK}} X,\ idb \xrightarrow{\text{FK}} X$$

stores the versions of a space. The *version*-attribute gives the version in which an object or bounded-by association appears for the first time, and VX is the object table of the *version space* described below. This yields a single space containing all objects appearing in some version. The following schema stores when (i.e. in which versions) an object or bounded-by association is deleted:

$$DelX(\underline{id}, version),\ id \xrightarrow{\text{FK}} X,\ DelR(\underline{ida}, \underline{idb}, version),\ (ida, idb) \xrightarrow{\text{FK}} R$$

The version numbers and transitions are stored as:

$$VX(\underline{version}),\ VR(\underline{fromv}, \underline{tov}),\ fromv \xrightarrow{\text{FK}} VX,\ tov \xrightarrow{\text{FK}} VX$$

which is the *version space* $V = (VX, VR)$, a topological data type for spaces of arbitrary combinatorial dimension. Since the version graph is a directed acyclic graph, the consistency rule for the version space is that of a T_0-space.

The *topological merge* of versions $(X_1, R_1), (X_2, R_2)$ which are each modifications of a common space (X, R) is a space (Y, S) fitting into the diagram:

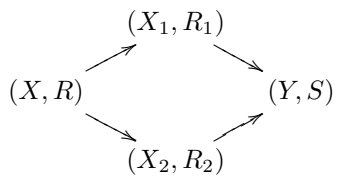

The set Y consists of all points of X occurring in X_1 or X_2, together with any additional points in either version. A *conflict* can occur if a point y is introduced in version X_1, and also in X_2. In that case, y will have the same *id*-value in both versions, and the conflict does occur if some attribute other than the *version*-value does not coincide in both versions. We call this an *inherent conflict*. In the same way, the topology-defining relation S is given as the set of all pairs from R occurring in R_1 or R_2, together with all new pairs from R_1 or R_2. Again, an inherent conflict occurs for pairs (a, b) appearing in both relations R_1 and R_2, if there exist attributes other than *version* and *id* which are different.

If there are no inherent conflicts, then the merge (Y, S) is a valid topological data type. However, there is a potential source for another type of conflict which we call *consistency conflict*. This happens if the topological space defined by (Y, S) violates some pre-assigned consistency rule. In the case of a conflict, a warning statement is issued, and it is up to the user to resolve the matter.

The merge of texts can be viewed as a topological merge by viewing a string as a linear DAG $\bullet \to \bullet \to \cdots \to \bullet$ whose semantic attribute takes values in an alphabet. If the resulting space is again a linear DAG without inherent conflict, then the topological merge of texts is valid. However, even without inherent conflict one can obtain topological spaces not representing text. E.g. the topological merge of texts in Fig. 4 exhibits this consistency conflict.

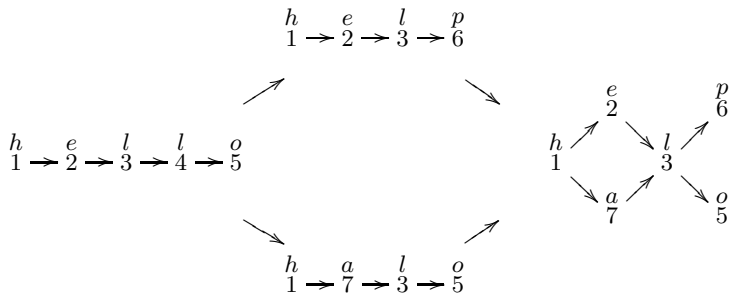

Fig. 4. A topological merge of texts violating the consistency rule "linear DAG".

Observe that the merge of two versions does not depend on the common source. The latter is only needed if the two versions are to be constructed from the modifications of the common source. But once the two spaces are known, it is only a matter of deciding which points and bounded-by associations they have in common, and where they supplement or possibly contradict each other, in order to obtain the topological merge.

As backtracking is the only way to produce existing versions, it follows that every version space is a T_0-space. We allow any finite T_0-space as a possible version space. In particular, a version space need not have a unique starting

point, or even be connected. This allows to begin with different parts of some spatial model. Each are successively modified, modifications are merged, and in the end one unique realisation is obtained through a final merge. In every step, a valid topological data type is obtained, and at each merge the occasional inherent and consistency conflict is resolved whenever it occurs.

To store a DAG as a topological data type comes natural. An alternative would be a one-dimensional simplicial complex. However, this increases the size by the relation which associates edges with boundary vertices, plus the additional orientation information of edges. Hence, increasing the dimension does in fact lead to a decrease in complexity.

Find all versions which were modified in order to obtain a given version. This queries for all elements x in the version space $V = (VX, VR)$ from which a directed path leads to given version v. These form the *minimal neighbourhood*

$$U_v := \{x \in VX \mid (x, v) \in VR^*\}, \tag{2}$$

where VR^* is the reflexive and transitive closure of VR. This set is obtained by the following a simple relational query which for convenience is given in SQL:

```
select fromv as version
from poVR
where poVR.tov = v;
```

where $poVR$, again, is the view on VR^* as described in Sec. 3. This minimal neighbourhood yields precisely the version numbers needed in order to build up the space belonging to v from the initial version(s).

Do given versions and their modifications reconstruct a given version? The question is whether all paths from the initial versions $V_0 \subset V$ to the given version $v \in V$ pass through the given set $W \subseteq V$ of versions. The answer is given by the following consideration: The minimal neighbourhood U_A of a set $A \subseteq V$ is

$$U_A := \{x \in VX \mid \exists a \in A\colon (x, a) \in VR^*\}, \tag{3}$$

and is the union of all paths ending in A. The paths going out of A are given by

$$\mathrm{cl}(A) := \{x \in VX \mid \exists a \in A\colon (a, x) \in VR^*\} \tag{4}$$

which is the *closure* of A in V. Hence, the paths from V_0 to v are given by

$$P(V_0, v) := U_v \cap \mathrm{cl}(V_0), \tag{5}$$

and the paths passing through W are all contained in $P(W) := U_W \cup \mathrm{cl}(W)$, from which we infer that the answer is given by the query:

$$\textit{Is it true that } P(v_0, v) \subseteq P(W)\textit{?} \tag{6}$$

This breaks down to queries for minimal neighbourhoods and closures, composed by union, intersection and inclusion queries. The relational query for the minimal neighbourhood of a set A is a slight modification of that for points:

```
select version from VX
where version in (
   select fromv
   from poVR
   where tov in (select id from A));
```

The closure cl(A) also uses the partial order $poVR$, and is achieved by reverses the order of pairs, hence simply by swapping the attributes *tov* and *fromv*, in the above SQL-statement. The relational query (6) can now be easily formulated, but we refrain from writing its precise statement.

6 Levels of Detail

It is often important to have spatial data at different levels of detail. In order to enable queries across different levels of detail it is necessary to link objects in one LoD with their aggregate object in the next LoD. As all objects in the finer LoD have a unique counterpart in the coarser LoD, we have a *generalisation function* $g\colon X \to Y$ from the fine space X to the coarse space Y. One important consistency rule is that objects that are "close" to each other must generalise to objects which are also "close" to each other. This "closeness" is a topological notion and determined by the bounded-by relation defining the topology of a finite space. Respecting that relation is nothing but to require that g be *continuous* [13, p. 508]. More precisely, g is a *continuous function* between spaces (X, R) and (Y, S) if every bounded-by association for X is mapped to a (possibly indirect) bounded-by association in Y: either $g(x_1) = g(x_2)$ or

$$(x_1, x_2) \in R \;\Rightarrow \exists\, y_1, \ldots, y_m \in Y\colon (g(x_1), y_1), \ldots, (y_m, g(x_2)) \in S$$

This rule is equivalent to the usual definition of continuous function from topology [18, §4]: namely that the pre-image of an open set be open [16, Ch. I §5]. Notice that the subspace topology from the last paragraph of Sec. 4 is defined in such a way that the inclusion function is continuous.

Another consistency rule is that g be surjective. Then every object in the coarse space is indeed the generalisation of an object in the fine space.

For the classical model

$$\mathit{Solid} \to \mathit{Face} \to \mathit{Edge} \to \mathit{Vertex} \tag{7}$$

a continuous function g for generalisation purposes implies the explicit modelling of up to $16 = 4^2$ possible types of association pairs, because g can map one class to any other class, and there is no reason to forbid the mapping of one class to a certain other class. Extending this to spaces of arbitrary dimension n (e.g. space-time etc.) gives $(n+1)^2$ different LoD associations to be modelled explicitly in order to form one single generalisation function. So, the classical model (7) considerably increases the complexity for describing functions. But if LoD associations are modelled on a common class of primitive objects, then the complexity of the class model decreases substantially. In the UML diagram of Fig.

5, this class is called `SpatialObject`. Instead of several dimension-dependent bounded-by associations there is simply one generic `BoundedBy` association between arbitrary objects of a finite space in the most general (topological) sense. Functions are incorporated as a `GeneralisesTo` association respecting the consistency rule: "continuous functions that are surjective between two consecutive LoDs". Of course, further consistency rules can be imposed if necessary.

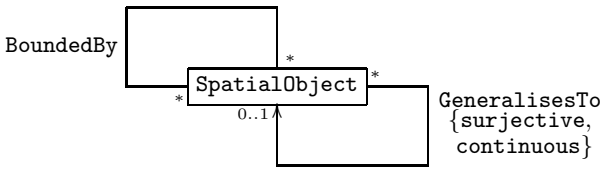

Fig. 5. The family of surjective and continuous `GeneralisesTo`-functions and the `BoundedBy`-relation on objects.

Assume that it is of interest if there is a path from a to b inside a given subset A of a space X with a generalisation function $g\colon X \to Y$. It would be economic to transport the question to Y and ask for a path from $g(a)$ to $g(b)$ inside the set $g(A)$, and then infer back to X. If A is connected, the answer is "yes". It is known that if A is connected, then $g(A)$ is also connected [14, Ch. 7, Thm. 1]. In case the pre-image of every connected subset of Y is connected, g is called *monotonic* [19, Ch. V, §46]. For monotonic generalisation functions it follows that if A is the full pre-image of $g(A)$, then the path query $a \rightsquigarrow b$ inside A can be reduced to the path query $g(a) \rightsquigarrow g(b)$ inside $g(A)$. In general, one has $A \subseteq g^{-1}(g(A))$, so that a positive answer for $g(a) \rightsquigarrow g(b)$ yields the existence of paths $a \rightsquigarrow b$ inside $g^{-1}(g(A))$ which need to be checked upon leaving A or not. Of course, by continuity, a negative answer for $g(a) \rightsquigarrow g(b)$ inside $g(A)$ implies a negative answer for $a \rightsquigarrow b$ inside A. The upshot is that the consistency rule "monotonic" allows to use g for accelerating the connectivity query for subsets A which are pre-images of generalised sets by delegating the query to the set $g(A)$ which in general is a smaller data set than the original A. Hence, the following *filtering* approach can be applied:

Algorithm 1 (Path query) Input. *Monotonic generalisation function* $g\colon X \to Y$*,* $A \subseteq X$*, and* $a, b \in A$*.*

Result. *"Yes" if* $a \rightsquigarrow b$ *in* A*, otherwise "No".*

Step 1. *Compute* $g(a), g(b), g(A)$*. Determine, if* $\exists\, g(a) \rightsquigarrow g(b)$ *inside* $g(A)$*.*

Step 2. *If "No", then answer* $=$ *No. Otherwise, determine if* $g^{-1}(g(A)) \subseteq A$*.*

Step 3. *If "Yes", then answer* $=$ *Yes. Otherwise, determine if* $\exists\, a \rightsquigarrow b$ *inside* A*.*

Step 4. *If "Yes", then answer* $=$ *Yes. Otherwise, answer* $=$ *No.*

Output. *answer.*

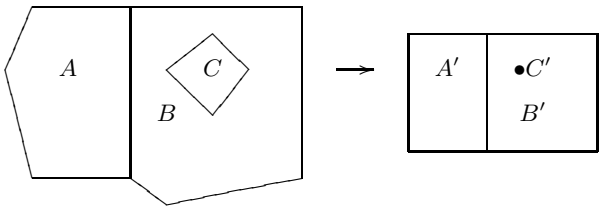

Fig. 6. A generalisation to a non-classical space.

Example. Assume the generalisation of a region subdivided into polygons A, B, C, as depicted in Fig. 6. There is a monotonic generalisation function to A', B', C'. Now polygon C generalises to vertex C' in the interior of polygon B'. However, the classical model (7) does not allow vertex C' to be in a direct bounded-by association with a face. Consequently, the information that there is a path from anywhere in A' to C' inside the generalised region cannot be inferred from the classical model without resorting to the underlying geometry. This has grave consequences. Namely, the geometry and topology of the generalised region are by force inconsistent: geometry says that C' lies inside the face B', but the classical model forces C' to be topologically disconnected from B'. Consequently, the generalisation function is not continuous! And since the statement: "B' is bounded by C'" cannot be explicitly modelled, it has to be inferred from the position of C'. But then, a possible position error of C' causing this vertex to be outside B' cannot be corrected by the topological model. But in that case, the incorrect geometry would be consistent with the classical topology model ...

This is remedied by topological data types. A direct bounded-by relation R associates B' with the four lines adjacent to B' which are themselves, through R, bounded by the four corners of B' which are not bounded by other objects. Hence, the latter are vertices, implying that B' must be a face. And R associates B' with C' which in turn is not bounded by an object. Hence, C' is a vertex at the boundary of object B'. And now the generalisation function is continuous, and furthermore monotonic. So, the question whether there is a path from anywhere in A to C can be delegated to the simpler generalised region.

We now show how interpolation between different LoDs is made possible. Assume that there is a chain $X_1 \to \cdots \to X_m$ of generalisation functions, and each space X_i is embedded in some $\mathbb{R}^n$. We assign to each X_i an extra level-coordinate $i \in \mathbb{R}$, leading to an embedding into $\mathbb{R}^{n+1} = \mathbb{R}^n \times \mathbb{R}$, and each LoD i sits inside some slice $\mathbb{R}^n \times \{i\}$. The generalisation function g_i associates each $x \in X_i$ with some $g_i(x) \in X_{i+1}$. If x and $g_i(x)$ are vertices, then they have, by our model, coordinates in $\mathbb{R}^{n+1}$, and an interpolating function between x and y can be defined. So, by assigning to *every* element x of every X_i coordinates (e.g. by taking a representative in the interior of the geometric realisation of x), it becomes possible to define interpolation functions between any object x and $g_i(x)$. The interpolation function can be given as a family of functions

$f_x\colon [0,1] \to \mathbb{R}^{n+1}$ with $f_x(0) = x$ and $f_x(1) = g_i(x)$, and as long as $t \in [0,1)$ the objects $f_x(t)$ are considered a placeholder for x inside X, but when $t = 1$ they become generalised to $f_x(1) = g_i(x)$. In this way, the topology of X_i stays unchanged while $t \in [0,1)$ and becomes that of X_{i+1} as soon as $t = 1$. The geometry can be made to gradually shrink from one LoD to the next. By giving coordinates to all elements, and not only the vertices, it becomes possible in a geometric realisation of a continuous zoom, for each vertex to have a unique trajectory, and so a unique position inside each slice between two LoDs. This includes also those vertices generalising to elements which are not vertices. In fact, unique trajectories become possible for all elements through the positions of their geometric representatives (i.e. coordinates) in $\mathbb{R}^{n+1}$.

From above, we derive the following relational schema for generalisations which allow interpolations between LoDs. Notice that the table *Vertex* from Sec. 3 is now called *Point*, as it may contain elements which are not vertices.

$$X(\underline{id}, \underline{lod}, gid, glod, atts), \quad (gid, glod) \overset{\text{CFK}(R)}{\longrightarrow} X$$
$$R(\underline{ida}, \underline{idb}, \underline{lod}), \quad (ida, lod) \overset{\text{FK}}{\longrightarrow} X,\ (idb, lod) \overset{\text{FK}}{\longrightarrow} X$$
$$Point(\underline{pid}, \underline{lod}, x, y, z, t), \quad (pid, lod) \overset{\text{FK}}{\longrightarrow} X$$

Here, $(gid, glod) \overset{\text{CFK}(R)}{\longrightarrow} X$ denotes a *continuous foreign key*, i.e. a foreign key which defines a continuous function from the set of all tuples in X having $(gid, glod) \neq (\texttt{NULL}, \texttt{NULL})$ to X with respect to the topology generated by R.

The attribute $R.lod$ in both foreign keys from R to X gives a disjoint union of an indexed family of spaces [16, Ch. I, §3]. This schema models each LoD in its entirety, whereas [7] propagates the idea that objects in one LoD not collapsing with other objects in the next LoD need not be repeated in the model. It would be interesting to compare both approaches.

7 Integrating the Different Spaces

Here, we put together the individual schemas for space-time, version and scale to one single schema. The different data models are integrated by collecting all tables, incorporating the attributes, and providing the foreign keys. As the LoD-attribute is part of the primary key for the point set, all elements (not only vertices) can be given space-time coordinates. This leads to the following tables:

$$X(\underline{id}, \underline{lod}, gid, glod, version, atts),\ R(\underline{ida}, \underline{idb}, \underline{lod}),\ Point(\underline{pid}, \underline{lod}, x, y, z, t)$$
$$DelX(\underline{id}, \underline{lod}, version),\ DelR(\underline{ida}, \underline{idb}, \underline{lod}, version),\ VX(\underline{version}),\ VR(\underline{fromv}, \underline{tov})$$

And the corresponding foreign keys are:

$$X.version \overset{\text{FK}}{\longrightarrow} VX, \quad (X.gid, X.glod) \overset{\text{CFK}(R)}{\longrightarrow} X$$
$$(R.ida, R.lod) \overset{\text{FK}}{\longrightarrow} X, \quad (R.idb, R.lod) \overset{\text{FK}}{\longrightarrow} X$$

$$(\mathit{Point.pid}, \mathit{Point.lod}) \xrightarrow{\text{FK}} X, \quad (\mathit{DelX.id}, \mathit{DelX.lod}) \xrightarrow{\text{FK}} X$$
$$(\mathit{DelR.ida}, \mathit{DelR.lod}) \xrightarrow{\text{FK}} X, \quad (\mathit{DelR.idb}, \mathit{DelR.lod}) \xrightarrow{\text{FK}} X,$$
$$\mathit{VR.fromv} \xrightarrow{\text{FK}} \mathit{VX}, \quad \mathit{VR.tov} \xrightarrow{\text{FK}} \mathit{VX}$$

All spaces made of polytopes in Euclidean space-time $\mathbb{R}^4$ can be consistently modelled with their versionings and generalisations. Although the combinatorial dimension can be arbitrary, the element coordinates are x, y, z, t. Introducing more coordinates increases the dimension of the embedding space $\mathbb{R}^n$. Other semantic data can be linked to the model by extending the database schema.

In which versions is there a path $a \rightsquigarrow b$ inside region A? This query across versions and LoDs is answered with the schema above. Assume that $a, b \in X$ and each point of $A \subseteq X$ are given by (*id*, *lod*). First, find all versions of A containing a and b. As elements enter X in the version given by *version*, and leave X in the version stored in *DelX*, the versions containing a and b are the intersection of the version intervals for the two points. These intervals are determined by the relation *VR* on the version table *VX*. For each corresponding *version*-value, there is a different version of A as a subspace of (X, R), determined from X, *DelX*, *DelR*, and for each version containing a and b, the query can be answered as follows. If a and b are in the same LoD, then its topology, stored in R, can be used. If a and b are in different LoDs, then the continuous foreign key CFK(R) is used for determining paths between those LoDs. If CFK(R) is monotonic, the query can be accelerated by Algorithm 1. This query uses the whole schema, except for *Point*. An explicit formulation in SQL is left to the reader.

The Integrated 6D+ Space. So far, we have a system of models for "elementary" spaces (cf. Sec. 2). To show how these can be combined into one space, we first extract the LoD space (*LoDX*, *LoDR*) by two queries that project *Xv* and *Rv*, a version v of the stored spaces, onto its two *lod* attributes:

```
create view LoDX as
   select lod from Xv union select glod as lod from Xv;
create view LoDR as select lod,glod from Xv;
```

A query converts this into the 1D edge graph (V, RV), V being the "union" of *LoDX* with *LoDR* after duplicating the identifiers x in *LoDX* into pairs $(\mathit{lod}, \mathit{glod}) = (x, x)$ to make it union compatible with *LoDR*. The relation *RV* contains $((a, b), (a, a))$ and $((a, b), (b, b))$ for every pair (a, b) in *LoDR*. Adding the graph *Gv* of the generalisation function in *Xv* to *Rv*, after making them union compatible, yields a new topology $\mathcal{T}(\mathit{Gv} \cup \mathit{Rv})$. By an *equi join* of spaces $(\mathit{Xv}, \mathit{Gv} \cup \mathit{Rv})$ and (V, RV) on *Xv.lod* and *V.lod*, we get an *equi-join space*, or a topological *pullback* [12, p. 406], of dimension ≥ 5. This is also a combinatorial variant of the *mapping telescope* [12, p. 312]. However, for each LoD i, except the coarsest, it contains two redundant copies of one space: the space at LoD vertex (x, x) and a homeomorphic copy at LoD edge $(x, g(x))$. So, the integrated 5D+ space has redundant information, which is not unusual with table joins.

One task of database design is to normalise by factoring redundant tables into smaller tables. Similarly, it is also be possible to integrate the whole versioning history and the LoDs, which gives an even bigger 6D+ space with even more redundancies. For the above schema this means: Whereas it is possible to formulate a query that integrates all spaces into one huge space, this space will have anomalies and thus may serve as a view for integrated space-time-version-view queries. However, it is not suitable as a relational schema for storage because of its anomalies. In short: We propose the future development of a topological relational database design theory extending its relational counterpart that has proved so successful in recent years.

8 Implementation

In Sec. 4, the subset Xt at a time point t of a set X in space-time was obtained by a relational selection. This was converted into a topological subspace. In fact, all basic query operators of relational algebra [8] can be turned into *topological relational query operators* operating on spaces and returning spaces, analog to their relational counterparts, as demonstrated in [9]. Here, we shortly describe the prototype on `pavel.gik.kit.edu`, currently designed as a first experimental implementation of the *semantics* of the query operators.

There are two classes of topological constructions [14, Ch. 3]: the initial (or "induced") spaces and the final (or "co-induced") spaces. A relational query operator on some input *spaces* operates on their elements and returns a *set* X. Now the result tuples are linked with the spatial entities from the input by functions: either from X back to the input (intersection, selection, or join), giving an *initial space*, or the function maps input entities to X (like union, projection), giving a *final space*.

The prototype is programmed in Common Lisp, and has its own simplified relational algebra in Lisp syntax. A space can be defined by the `space` constructor whose input is a set, the topology-defining relation and the two foreign keys. Each basic operator `op` has its spatial counterpart `op-space`, like `natjoin-space`, or `project-space`, that acts on the sets, constructs the corresponding (initial or final) topology for the result set and returns a space. These operators can be arbitrarily nested. The experiences gained will help to produce a *topological relational database management system* which should provide the topological data modelling presented above as builtin feature. It would also provide topological consistency rules, and could be a starting point for the discussion of topological data modelling rules towards a topological data modelling theory extending the current relational modelling theory.

9 Conclusion

A relational database schema based on Alexandrov topology, that seamlessly integrates 4D space-time data, version histories, and different levels of detail (LoD), is presented. Such a topology can always be represented by a directed

acyclic graph, and it imposes fewer restrictions than the canonical *Solid-Face-Edge-Vertex*-model for spatial data. The gained flexibility and simplicity alleviates more sophisticated spatial data modelling, and endows spatial data with an Alexandrov topology, which has practical consequences: As topology is fundamental it is likely to have more to offer for spatial data modelling than is momentarily used. A first contribution is a precise definition of "spatial data dimension". A topological version space allows the recovery of different versions of a spatial model by using queries based on topological constructions. Among the new consistency rules, "continuity" of foreign keys allows consistent modelling in different LoDs, and "monotonicity" allows accelerated path queries. Also, topological queries across versions and LoDs are enabled.

Acknowledgements.

This work is funded by the DFG with grants BR2128/12-2 and BR3513/3-2.

References

1. Sakr, M.A., Güting, R.H.: Spatiotemporal pattern queries in secondo. In Mamoulis, N., et al., eds.: SSTD 2009, LNCS 5644. (2009) 422–426
2. Anh, N., Vinh, P.T., Duy, H.: A study on 4d gis spatio-temporal data model. In: Knowledge and Systems Engineering (KSE), 2012 Fourth Int. Conf. (2012) 34 –38
3. van Oosterom, P., Stoter, J.: Principles of 5d modeling, full integration of 3d space, time and scale. Geospatial World Weekly (2012)
4. Lienhardt, P.: n-dimensional generalized combinatorial maps and cellular quasi-manifolds. Int. J. Comp. Geom. App. **4** (1994) 275–32
5. Engelking, R.: Dimension theory. North-Holland (1978)
6. Hartshorne, R.: Algebraic Geometry. Springer (1993)
7. van Oosterom, P., Meijers, M.: Towards a true vario-scale structure supporting smooth-zoom. In: Proc. 14th Workshop ICA Commission on Generalisation and Multiple Representation, Paris (2011)
8. Codd, E.F.: The Relational Model for Database Management: Version 2. Addison-Wesley (1990)
9. Bradley, P.E., Paul, N.: Using the relational model to capture topological information of spaces. The Computer Journal **53** (2010) 69–89
10. Zlatanova, S., Rahman, A.A., Shi, W.: Topological models and frameworks for 3d spatial objects. Computers & Geosciences **30**(4) (2004) 419 – 428
11. Bradley, P.E., Paul, N.: Dimension of alexandrov topologies. Technical Note (2013) http://pavel.gik.kit.edu/doc/dimension-alex.pdf.
12. Hatcher, A.: Algebraic Topology. Cambridge University Press (2002) online: `http://www.math.cornell.edu/~hatcher/`.
13. Alexandroff, P.: Diskrete Räume. Matematiećeskij Sbornik **44**(2) (1937) 501–519
14. Adamson, I.T.: A General Topology Workbook. Birkhäuser (1995)
15. Abbott, E.A.: Flatland: A romance of many dimensions. Seeley, London (1884)
16. Jänich, K.: Topology. Springer (1995)
17. Codd, E.F.: Extending the Database Relational Model to Capture More Meaning. ACM Trans. Database Syst. **4**(4) (1979) 397–434
18. McCord, M.C.: Singular homology and homotopy groups of finite topological spaces. Duke Math. J. **33** (1966) 465–474
19. Kuratowski, K.: Topology. Volume II. Acad. Press (1968)